

SIMULATION OF HARDWARE AND SOFTWARE

TECHNICAL FIELD OF THE INVENTION

5 This invention relates to the field of digital system design verification. In particular, the invention relates to optimising the speed of simulation of a digital system design that features both software and hardware components, providing a method of allowing a C
10 model to control the simulation in the HDL environment and allowing the simulation of interactive programs.

BACKGROUND OF THE INVENTION

15 A large number of the digital systems being designed today are task-specific systems that consist of standard or custom hardware and standard or custom software.

20 Some systems have a predominately software element; in that most of the design effort is spent implementing the system in software. Typically, this software is implemented on standard or previously designed hardware.

25 Other systems are hardware dominant; in that most of the design effort is spent implementing the system in programmable logic devices (PLDs) or application specific integrated circuits (ASICs).

30 In either case where hardware or software dominates, there are a variety of tools that can be used by the designer to validate the system and ensure that it is operating correctly. In the case of a software
35 dominant system, debugger and other related software tools could be used. In the case of a hardware

dominant system, hardware emulation or simulation could be used.

5 The systems that are the most difficult to validate are those that are neither hardware nor software dominant, where both parts play an important role in the success of the system. Often, the hardware and software are initially developed separately with each part being tested independently of the other. The hardware is
10 tested using a hardware simulator or emulator. The software is validated using an instruction set simulator on a computer. Typically, an instruction set simulator will run at speeds that are several orders of magnitude higher than in a hardware simulator.

15 Once both the hardware and software have been verified separately, the two are co-simulated. That is, a simulation is carried out where the software runs on a simulation of the prototype hardware.

20 Due to the much lower speed of hardware simulation, this co-simulation is considerably slower than instruction set simulation.

25 **SUMMARY OF THE INVENTION**

A method is provided that provides fast software simulation within a hardware description language (HDL)
30 simulation environment.

According to a first aspect of the present invention, a method of co-verifying a hardware and software system is provided that comprises simulating an embedded

system (including any processors) in a separate thread within a HDL simulation environment.

The embedded system is implemented in a C model that runs asynchronously to the HDL design, which allows the software executing on the embedded processor to run ahead of the HDL simulation. This allows fast simulation of processor instructions within the embedded system model.

10

According to an embodiment of the present invention, the C model is synchronised with the HDL simulation via a reference clock, which limits the number of microprocessor clock periods executed per reference clock period.

15

According to a second aspect of the present invention, a method of allowing the C model to control the HDL simulation is provided that comprises: running a macro in the HDL simulation to set up a TCP/IP socket, connecting the C model to the socket, allowing the C model to control the HDL simulation.

20

According to a third aspect of the present invention, a method for providing an IO interface to a simulation model to allow the simulation of interactive programs is provided that comprises: simulating an embedded input/output device within the simulation model; connecting the model of the embedded input/output device to a terminal emulator over TCP/IP; and running a program that interacts with the user via the terminal emulator.

30

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows the interface between a C model and the HDL simulation in accordance with a first aspect of the invention.

Figure 2 shows the use of multiple threads in simulating the hardware components in accordance with a first aspect of the invention.

Figure 3 shows the execution of the co-verification system running on a host machine with a single processor.

Figure 4 shows the execution of the co-verification system running on a host machine with two processors.

Figure 5 shows the simulation of a C model thread and a HDL simulation thread asynchronously using the thread scheduling of the host machine.

Figure 6 illustrates a software debugger interfacing with a HDL simulation environment over a TCP/IP connection.

Figure 7 shows conventional intra-process communication between a C model and a HDL simulation.

Figure 8 illustrates a software debugger interfacing with a C model in accordance with the second aspect of the present invention.

Figure 9 illustrates a software debugger interfacing directly with the HDL simulation in accordance with the second aspect of the present invention.

Figure 10 illustrates an exemplary method for providing IO interface to a simulation model to allow the simulation of interactive programs according to a third aspect of the present invention.

5

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

In accordance with a first aspect of the invention, fast simulation of software within a hardware
10 description language (HDL) simulation environment is achieved by implementing any tightly coupled processors and peripheral devices within a single C/C++ model. Bus functional models are used to provide an interface between the C/C++ model and more loosely coupled
15 hardware components simulated in HDL.

Although the embodiment of the invention described below refers to embedded processors and embedded peripherals, it will be appreciated that the invention
20 is also applicable to processors and peripherals that are not embedded. It will also be appreciated that the references to HDL and C/C++ models are not to be considered restrictive. Other programming languages, such as Java, assembly language, or Visual Basic could
25 be used in accordance with the invention. Modelling techniques that are covered by the term "HDL" include all hardware description languages, such as Verilog™ and VHDL. However, the present invention should not be limited to these examples of programming language and
30 hardware description language.

The C model is implemented such that both the embedded processor and all embedded peripherals to the embedded processor are modelled in C/C++. Implementing the
35 tightly coupled aspects of the system (i.e., processor,

any peripherals with which the processor interacts and internal memory interfaces) in C/C++ allows the embedded processor and peripherals to interact within a single C/C++ model, which means that communication
 5 between the C model and the HDL model is reduced, dramatically improving performance over an implementation which contains multiple interfaces between blocks of HDL and C/C++.

10 Figure 1 shows an example of the interface between a C model and the HDL simulation in accordance with the first aspect of the invention. As described above, the embedded processor and peripherals 2 are implemented in C/C++ code, the PLD design 4, PLD application interface
 15 6 and the bus functional model 8 are modelled in a hardware description language such as Verilog or VHDL.

The two simulation models communicate over a programming language interface (PLI) or a foreign
 20 language interface (FLI) 10.

It should be noted that although only a single PLD design is illustrated in Figure 1, any number of PLD/ASICs could be modelled in the HDL simulation.

25 The use of a bus functional model 8 has been extended so that it is driven by the C model of the embedded processors and peripherals, instead of script files, as is conventional. A C/C++ model-to-PLD bridge, that
 30 synchronises transactions across clock domains, is implemented such that the slave is modelled in C/C++ code with the master modelled in HDL. The situation is reversed for the PLD-to-C/C++ model bridge, with the slave in HDL and the master implemented in C/C++.

35

Further in accordance with the first aspect of the present invention, the C/C++ model of the embedded processor and peripherals runs asynchronously to the HDL simulation.

5

Figure 2 shows the use of multiple threads in simulating the hardware components of the system. The C/C++ model 20 is in a separate thread of execution on the host machine to the HDL simulation 22. The C/C++
 10 model thread 20 is spawned within a PLI/FLI function 24.

Synchronisation between the C model 20 and the HDL simulation 22 can be achieved through an HDL parameter
 15 that specifies a maximum number of processor clock cycles simulated in the C/C++ model 20 per reference clock cycle.

Figure 3 shows the execution of the co-verification system running on a host machine with a single
 20 processor. With the parameter set to 4, the C/C++ model is advanced by 4 processor clock periods for each reference clock period in the HDL simulation thread.

25 Shaded blocks 30 in Figure 3 represent a single processor clock period in the C/C++ model thread while the unshaded blocks 32 represent a single reference clock cycle in the HDL simulation thread.

30 Figure 4 shows the execution of the co-verification system running on a host machine with two processors. The first processor runs the C model and the second processor runs the HDL simulation. Again, with the parameter set to 4, the C model is advanced by 4
 35 processor clock periods on the first processor, before

the HDL simulation thread is advanced by a single reference clock period on the second processor.

Shaded blocks 40 in Figure 4 represent a single processor clock period in the C/C++ model thread while the unshaded blocks 42 represent a single reference clock cycle in the HDL simulation thread.

In a multi-processor host machine, a mode is provided, selectable through an HDL parameter, which simulates the C/C++ model thread and the HDL simulation thread asynchronously using the thread scheduling of the host machine. This mode is illustrated in Figure 5.

Typically, in this mode, hundreds to thousands of processor clock cycles in the C/C++ model thread are executed per reference clock cycle.

To avoid the insertion of spurious wait cycles when C/C++ model-to-PLD interactions occur, the threads are synchronised around such interactions. In this case, the performance gain of the system is balanced against a loss of repeatability of simulation.

It should be noted that since the C/C++ model thread contains both the processor and peripherals, and combined with the multi-threaded nature of the implementation and the thread synchronisation, this provides a suitable environment to allow code running on the model of the embedded processor in the C/C++ model to run ahead of the HDL simulation very quickly. This is particularly advantageous when executing large fragments of software on the embedded processor without interaction with the PLD, such as running an Operating System application.

The implementation described allows the simulation of over three hundred thousand instructions per second within a HDL simulation environment on a standard home computer.

5

According to the second aspect of the present invention a software/hardware co-simulation environment is presented that provides full debug control (watchpoint, breakpoint, single stepping, etc.) over the software simulated on the embedded processor in the C/C++ model, whilst remaining in an HDL simulation environment. A high level of visibility of the C/C++ model-to-PLD interactions and/or software/hardware interactions is achieved.

15

Figure 6 shows a software debugger interfacing with the HDL simulation environment over a TCP/IP (Transmission Control Protocol/Internet Protocol) connection. It will be appreciated that any inter-process communication protocol, such as pipes, or file-based transfer, could be used in place of TCP/IP. The C/C++ model 60 and the HDL simulation 61 are implemented in the same process 62 so that communication between the HDL simulation 61 and the C/C++ model 60 can be most efficient. Such communication can be (but is not limited to) bus transactions, IO pin changes or interrupt signals.

20

25

Communications between the software debugger 64 and the C/C++ model 60 (represented by arrow 65) will generally be less host processor intensive than any HDL-to-C/C++ model communications (represented by the dotted arrow). Therefore, the software debugger 64 and C/C++ model 60 communications are more suitable to be implemented using inter-process communication protocols, such as TCP/IP.

30

35

Conventionally, any HDL simulator to C/C++ model communication, whether implemented over a programming language interface (PLI) or foreign language interface (FLI), results in the C/C++ model effectively being a slave to the HDL simulator. This is illustrated in Figure 7. The HDL simulation 70 sends control commands to the C/C++ model 72, to which the C/C++ model 72 responds. This means that the HDL simulator governs control of the C/C++ model; the C/C++ model is only able to influence execution using control requests as a response. Consequently, if the HDL simulation is paused - through a user or C/C++ model request - only the HDL simulator can resume overall simulation.

Figure 8 illustrates the system according to the second aspect of the present invention. The C/C++ model 80 is implemented within the same process as the HDL simulation 82. The HDL simulation 82 can send control commands to the C/C++ model 80 using an intra-process protocol (represented by the dashed arrows). The C/C++ model 80 is able to respond using the same protocol. A software debugger 84 and the C/C++ model 80 communicate over an inter-process communication protocol (represented by the solid arrows), such as TCP/IP. The software debugger 84 can send control commands to the C/C++ model 80, to which the C/C++ model 80 can respond.

With the software debugger 84 attached to the C/C++ model 80, the software debugger 84 has no knowledge of whether the C/C++ model 80 is attached to the HDL simulation 82 (and does not need to have such knowledge). Control of the C/C++ model 80 from the software debugger 84 - in terms of running, single-steps etc. - requires the C/C++ model 80 to have

control over the HDL simulation 82. As outlined above, this control cannot be achieved over a FLI or PLI.

Therefore, in accordance with the invention, the C/C++
5 model 80 communicates with a HDL simulation user
interface/front end 86 using an inter-process
communication protocol, such as TCP/IP. This allows
the C/C++ model 80 to send control commands to the HDL
simulation user interface 86. The HDL simulation front
10 end 86 is used to run a macro, setting up a listening
TCP/IP socket. When simulating an HDL design, the
C/C++ model 80 connects to this socket to allow it to
start/resume the HDL simulation 82. Thus, when the
software debugger 84 issues a start command, the C/C++
15 model 80 uses the TCP/IP link to the HDL simulation
front end 86 in order to resume HDL simulation 82.
When a stop command is issued from the software
debugger 84, the C/C++ model 80 can request simulation
stop via the normal intra-process communication.

20
Figure 9 illustrates an alternative system according to
the second aspect of the present invention. Again, the
C/C++ model 80 is implemented within the same process
as the HDL simulation 82. The HDL simulation 82 can
25 send control commands to the C/C++ model 80 using an
intra-process protocol (represented by the dashed
arrows). The C/C++ model 80 is able to respond using
the same protocol.

30 In accordance with the invention, the software debugger
84 communicates with a HDL simulation user
interface/front end 86 using an inter-process
communication protocol, such as TCP/IP. This allows
the software debugger 84 to send control commands to
35 the HDL simulation user interface 86. The HDL
simulation front end 86 is used to run a macro, setting

up a listening TCP/IP socket. When simulating an HDL design, the software debugger 84 connects to this socket to allow it to start/resume the HDL simulation 82. Thus, when the software debugger 84 issues a start
5 command over the TCP/IP link to the HDL simulation front end 86, the HDL simulation 82 is resumed.

The software debugger 84 can also indirectly control the C/C++ model 80 through communications over TCP/IP
10 to the HDL simulation front end 86, and then from the HDL simulation 82 to the C/C++ model 80 using an intra-process communication protocol.

According to the third aspect of the present invention,
15 a method for providing an IO interface to a simulation model to allow the simulation of interactive programs is provided.

As shown in Figure 10, an IO based peripheral device,
20 such as an embedded UART (universal asynchronous receiver-transmitter), is connected within the simulation model to a terminal emulator over TCP/IP or other inter-process communications protocol or technique (e.g., pipes, file transfer). Figure 10
25 illustrates the connection of an input/output peripheral device to a terminal emulator using an inter-process communications link, such as the preferred TCP/IP link. The C/C++ model processing thread communicates with the I/O device processing
30 thread within the simulation model.

The I/O device is connected to a terminal emulator, such as telnet, using the TCP/IP link. The terminal emulator is used to receive character, and other
35 instruction, inputs from a user. The user thereby makes use of a familiar terminal emulator to provide

control over the hardware simulation. By specifying appropriate settings (such as which terminal emulator and which port) the model can either (a) automatically spawn and connect to the terminal emulator or (b) wait
5 for a connection from the terminal emulator over TCP/IP before beginning simulation. This can provide an IO interface to the model through the UART that allows the simulation of interactive programs.

- 10 To allow fast and concurrent input and output, via the terminal emulator, a separate UART thread of execution is used to receive user input characters.

This method can readily be applied to other IO based
15 peripherals such as Ethernet MACs.

There is thus provided a method and apparatus for providing fast software simulation within a hardware description language (HDL) environment.

20

While the present invention has been described herein with reference to particular embodiments, various changes, and substitutions are possible within the spirit of the present invention. The invention is not
25 be limited to the particular embodiment disclosed, but rather it is defined by the scope of the claims that follow.